

# Fast Approximate Bayesian Computation for discretely observed Markov models using a factorised posterior distribution

Simon R. White

MRC Biostatistics Unit, Cambridge

`simon.white@mrc-bsu.cam.ac.uk`

Theodore Kypraios

University of Nottingham

`theodore.kypraios@nottingham.ac.uk`

Simon P. Preston

University of Nottingham

`Simon.Preston@nottingham.ac.uk`

December, 17, 2012

December, 17, 2012

## Abstract

Many modern statistical applications involve inference for complicated stochastic models for which the likelihood function is difficult or even impossible to calculate, and hence conventional likelihood-based inferential techniques cannot be used. In such settings, Bayesian inference can be performed using Approximate Bayesian Computation (ABC). However, in spite of many recent developments to ABC methodology, in many applications the computational cost of ABC necessitates the choice of summary statistics and tolerances that can potentially severely bias the estimate of the posterior.

We propose a new “piecewise” ABC approach suit-

able for discretely observed Markov models that involves writing the posterior density of the parameters as a product of factors, each a function of only a subset of the data, and then using ABC within each factor. The approach has the advantage of side-stepping the need to choose a summary statistic and it enables a stringent tolerance to be set, making the posterior “less approximate”. We investigate two methods for estimating the posterior density based on ABC samples for each of the factors: the first is to use a Gaussian approximation for each factor, and the second is to use a kernel density estimate. Both methods have their merits. The Gaussian approximation is simple, fast, and probably adequate for many applications. On the other hand, using instead a kernel density estimate has the benefit of consistently estimating the true ABC posterior as the number of ABC samples

tends to infinity. We illustrate the piecewise ABC approach for three examples; in each case, the approach enables “exact matching” between simulations and data and offers fast and accurate inference.

## 1 Introduction

Stochastic models are commonly used to model processes in the physical sciences [Wilkinson, 2011, Van Kampen, 2007]. For many such models the likelihood is difficult or costly to compute making it infeasible to use conventional inference techniques such as maximum likelihood estimation. However, provided it is possible to simulate from a model, then “implicit” methods such as Approximate Bayesian Computation (ABC) methods enable inference without having to calculate the likelihood. These methods were originally developed for applications in population genetics [Pritchard et al., 1999] and human demography [Beaumont et al., 2002], but are now being used in a wide range of fields including epidemiology [McKinley et al., 2009], evolution of species [Toni et al., 2009], finance [Dean et al., 2011], and evolution of pathogens [Gabriel et al., 2010], to name a few.

Intuitively, ABC methods involve simulating data from the model using various parameter values and making inference based on which parameter values produced realisations that are “close” to the observed data. Suppose that our data are a set of observations denoted  $\mathcal{X} = \{x_1, \dots, x_n\} \equiv \{x(t_1), \dots, x(t_n)\}$  of state variable  $X(t) \in \mathbb{R}^m$  at time points  $t_1, \dots, t_n$ . We assume that the data arise from a Markov stochastic model (this encompasses IID data as a special case) parameterised by the vector  $\theta$ , which is the target of inference. Prior beliefs about  $\theta$  are

expressed via the density  $\pi(\theta)$ . The following Algorithm 1 generates *exact* samples from the Bayesian posterior density  $\pi(\theta|\mathcal{X})$  which is proportional to  $\pi(\mathcal{X}|\theta)\pi(\theta)$ :

---

### Algorithm 1

Exact Bayesian Computation (EBC)

---

- 1: Sample  $\theta^*$  from  $\pi(\theta)$ .
  - 2: Simulate dataset  $\mathcal{X}^*$  from the model using parameters  $\theta^*$ .
  - 3: Accept  $\theta^*$  if  $\mathcal{X}^* = \mathcal{X}$ , otherwise reject.
  - 4: Repeat.
- 

This algorithm is only of practical use if  $X(t)$  is discrete, else the acceptance probability in Step 3 is zero. For continuous distributions, or discrete ones in which the acceptance probability in step 3 is unacceptably low, Pritchard et al. [1999] suggested the following algorithm:

---

### Algorithm 2

Approximate Bayesian Computation (ABC)

---

As Algorithm 1, but with step 3 replaced by:

- 3': Accept  $\theta^*$  if  $d(s(\mathcal{X}), s(\mathcal{X}^*)) \leq \varepsilon$ , otherwise reject.
- 

where  $d(\cdot, \cdot)$  is a distance function, usually taken to be the  $L_2$ -norm of the difference between its arguments;  $s(\cdot)$  is a function of the data; and  $\varepsilon$  is a tolerance. Note that  $s(\cdot)$  can be the identity function but in practise, to give tolerable acceptance rate, it is usually taken to be a lower dimensional vector of summary statistics that characterise key aspects of the data.

The output of the ABC algorithm is a sample from the ABC posterior density  $\tilde{\pi}(\theta|\mathcal{X}) =$

$\pi(\theta|d(s(\mathcal{X}), s(\mathcal{X}^*)) \leq \varepsilon)$ . Provided  $s(\cdot)$  is sufficient for  $\theta$ , then the ABC posterior density converges to  $\pi(\theta|\mathcal{X})$  as  $\varepsilon \rightarrow 0$ . However, in practise it is rarely possible to use an  $s(\cdot)$  which is sufficient, or to take  $\varepsilon$  especially small (or zero). Hence ABC requires a careful choice of  $s(\cdot)$  and  $\varepsilon$  to try to make the acceptance rate tolerably large, at the same time as trying not to make the ABC posterior too different from the true posterior,  $\pi(\theta|\mathcal{X})$ . In other words, there is a balance which involves trading off Monte Carlo error with “ABC error” owing to the choice of  $s(\cdot)$  and tolerance  $\varepsilon$ .

Over the last decade, a wide range of extensions to the original ABC algorithm have been developed, including Markov Chain Monte Carlo (MCMC) [Marjoram et al., 2003] and Sequential Monte Carlo (SMC) [Toni et al., 2009] implementations, the incorporation of auxiliary regression models [Beaumont et al., 2002, Blum and François, 2010], and (semi-)automatic choice of summary statistics [Fearnhead and Prangle, 2012]; see Marin et al. [2011] for a review. In all of these ABC variants, however, computational cost remains a central issue, since it is the computational cost that determines the balance that can be made between controlling Monte Carlo error and controlling bias arising from using summary statistics and/or non-zero tolerance.

In this paper we propose a novel algorithm called *piecewise ABC* (PW-ABC), the aim of which is to substantially reduce the computational cost of ABC. The algorithm is applicable to a particular (but fairly broad) class of models, namely those with the Markov property and for which the state variable is observable at discrete time points. The algorithm is based on a factorisation of the posterior density such that

each factor corresponds to only a subset of the data. The idea is to apply the ABC algorithm for each factor (a task which is computationally very cheap), to compute the density estimates for each factor, and then to estimate the full posterior density as the product of these factors. Taking advantage of the factorisation lowers the computational burden of ABC such that the choice of summary statistic and tolerance—and the accompanying biases—can potentially be avoided completely.

In the following section we describe PW-ABC in more detail. The main question of the method is how to use the ABC samples from each posterior factor to estimate the full posterior density. We discuss two approaches to estimating the relevant densities and products of densities, then we apply PW-ABC, using both approaches, to three examples: a toy illustrative example of inferring the probability of success in a binomial experiment, an autoregressive time-series model, and a dynamical predator–prey model. We conclude with a discussion of the strengths and limitations of PW-ABC, and of potential further generalisations.

## 2 Piece-wise ABC (PW-ABC)

Our starting point is to use the Markov property to write the likelihood as

$$\begin{aligned} \pi(\mathcal{X}|\theta) &= \left( \prod_{i=2}^n \pi(x_i|x_{i-1}, \dots, x_1, \theta) \right) \pi(x_1|\theta) \\ &= \left( \prod_{i=2}^n \pi(x_i|x_{i-1}, \theta) \right) \pi(x_1|\theta). \end{aligned} \quad (1)$$

The likelihood contribution of the first observation  $x_1$  is asymptotically irrelevant as the number of observations,  $n$ , increases and, henceforth, to keep the

presentation simple, we ignore the term  $\pi(x_1|\theta)$  in (1). Accounting for this, and by using multiple applications of Bayes' theorem, the posterior density can be written in the following factorised form,

$$\begin{aligned}\pi(\theta|\mathcal{X}) &\propto \pi(\mathcal{X}|\theta)\pi(\theta) \\ &= \left( \prod_{i=2}^n \frac{\pi(x_i|x_{i-1}, \theta)\pi(\theta)}{\pi(\theta)} \right) \pi(\theta) \\ &\propto \pi(\theta)^{(2-n)} \left( \prod_{i=2}^n \varphi_i(\theta) \right),\end{aligned}\quad (2)$$

where

$$\begin{aligned}\varphi_i(\theta) &= c_i^{-1} \pi(x_i|x_{i-1}, \theta)\pi(\theta) \\ c_i &= \int \pi(x_i|x_{i-1}, \theta)\pi(\theta)d\theta.\end{aligned}$$

Essentially, in (2) the posterior density,  $\pi(\theta|\mathcal{X})$ , of  $\theta$  given the full data  $\mathcal{X}$  has been decomposed into a product involving densities  $\varphi_i(\theta)$ , each of which depends only on a pair of data points,  $\{x_{i-1}, x_i\}$ .

The key idea now is to use ABC to draw approximate samples from each of the densities  $\varphi_i(\theta)$ . Applying Algorithm 2 involves (i) drawing  $\theta^*$  from  $\pi(\theta)$ , (ii) simulating  $x_i^*|x_{i-1}, \theta^*$ , and (iii) accepting  $\theta^*$  if  $d(s(x_i), s(x_i^*)) \leq \varepsilon$ . We use  $\tilde{\varphi}_i(\theta)$  to denote the implied ABC density from which these samples are drawn (with  $\tilde{\varphi}_i(\theta) = \varphi_i(\theta)$  if  $s(\cdot) = \text{Identity}(\cdot)$  and  $\varepsilon = 0$ ). By repeating (i)–(iii) we generate samples of, say,  $m$  draws,  $\theta_{i(1)}^*, \dots, \theta_{i(m)}^*$ , from each  $\tilde{\varphi}_i(\theta)$ . Now, suppose that  $\hat{\varphi}_i(\theta)$  is an estimate, based on the sample  $\theta_{i(1)}^*, \dots, \theta_{i(m)}^*$ , of the density  $\tilde{\varphi}_i(\theta)$  (and hence of the density  $\varphi_i(\theta)$ ). Then the posterior density (2) can be estimated by

$$\hat{\pi}(\theta|\mathcal{X}) = g(\theta) / \int g(\theta) d\theta, \quad (3)$$

where

$$g(\theta) = \pi(\theta)^{(2-n)} \left( \prod_{i=2}^n \hat{\varphi}_i(\theta) \right). \quad (4)$$

The steps of PW-ABC are summarised in Algorithm 3.

---

**Algorithm 3** Piece-Wise Approximate Bayesian Computation (PW-ABC)

---

**for**  $i = 2$  to  $n$  **do**

- a: Apply the ABC Algorithm to draw  $m$  approximate (or exact, if  $s(\cdot) = \text{Identity}(\cdot)$  and  $\varepsilon = 0$ ) samples,  $\theta_{i(1)}^*, \dots, \theta_{i(m)}^*$ , from  $\varphi_i(\theta)$ ;
- b: Using the samples  $\theta_{i(1)}^*, \dots, \theta_{i(m)}^*$  and either (6) or (12), calculate a density estimate,  $\hat{\varphi}_i(\theta)$ , of  $\varphi_i(\theta)$ .

**end for**

Substitute the density estimates  $\hat{\varphi}_i(\theta)$  into (3) to calculate an estimate,  $\hat{\pi}(\theta|\mathcal{X})$ , of  $\pi(\theta|\mathcal{X})$ .

---

On the question of how to calculate the density estimates,  $\hat{\varphi}_i(\theta)$ , below we discuss two approaches: (i) using a Gaussian approximation, and (ii) using a kernel density estimate. Henceforth, quantities based on (i) are denoted by superscript g, and those based on (ii) are denoted by superscript k. In both cases we discuss the behaviour of the estimators in the asymptotic regime in which the number of observations,  $n$ , is kept fixed while the size of each ABC sample increases,  $m \rightarrow \infty$ .

## 2.1 Gaussian approximation for $\hat{\varphi}_i(\theta)$

Denote the  $d$ -dimensional multivariate Gaussian density with mean,  $\mu$ , and covariance,  $\Sigma$ , by

$$K(\theta; \mu, \Sigma) = (2\pi)^{-d/2} (\det \Sigma)^{-1/2} \exp \left( -\frac{1}{2} (\theta - \mu)^T \Sigma^{-1} (\theta - \mu) \right). \quad (5)$$

A Gaussian approximation for  $\hat{\varphi}_i(\theta)$  is

$$\hat{\varphi}_i^g(\theta) = K(\theta; \bar{\theta}_i^*, Q_i), \quad (6)$$

where

$$\begin{aligned} \bar{\theta}_i^* &= \frac{1}{m} \sum_{j=1}^m \theta_{i(j)}^*, \\ Q_i &= \frac{1}{m-1} \sum_{j=1}^m (\theta_{i(j)}^* - \bar{\theta}_i^*)(\theta_{i(j)}^* - \bar{\theta}_i^*)^T, \end{aligned}$$

are the sample mean and sample covariance of the ABC posterior sample  $\theta_{i(1)}^*, \dots, \theta_{i(m)}^*$ . A consequence of using (6) is that the product of the density approximations is also Gaussian:

$$\prod_{i=2}^n \hat{\varphi}_i^g(\theta) = w \cdot K(\theta; a, B), \quad (7)$$

where

$$B = \left( \sum_{i=2}^n Q_i^{-1} \right)^{-1}, \quad (8)$$

$$a = B \left( \sum_{i=2}^n Q_i^{-1} \bar{\theta}_i^* \right), \quad (9)$$

$$w = \det(2\pi B)^{1/2} \prod_{i=2}^n \det(2\pi Q_i)^{-1/2} \times \prod_{s=2}^n \prod_{t>s}^n \exp \left( -\frac{1}{2} (\bar{\theta}_s^* - \bar{\theta}_t^*)^T R_{st} (\bar{\theta}_s^* - \bar{\theta}_t^*) \right), \quad (10)$$

$$R_{st} = Q_s^{-1} B Q_t^{-1}. \quad (11)$$

We note the following properties of approximation (6). If the densities  $\hat{\varphi}_i(\theta)$  from which the  $\theta_{i(1)}^*, \dots, \theta_{i(m)}^*$  are drawn are Gaussian, i.e.,  $\tilde{\varphi}_i(\theta) = K(\theta; \mu_i, \Sigma_i)$ , then  $\bar{\theta}_i^*$  and  $Q_i$  are unbiased and consistent estimators of  $\mu_i$  and  $\Sigma_i$ , respectively, and hence  $a$  and  $B$  are consistent estimators of the true mean and covariance of  $\prod \tilde{\varphi}_i(\theta)$ . More generally, for  $\tilde{\varphi}_i(\theta)$  which is not necessarily Gaussian,  $\bar{\theta}_i^*$  and  $Q_i$  are

consistent estimators of the mean and the variance of the Gaussian density,  $\hat{\varphi}_i^g(\theta)$ , which minimises the Kullback–Leibler divergence,

$$\text{KL}(\tilde{\varphi}_i(\theta) \parallel \hat{\varphi}_i^g(\theta)) = \int \tilde{\varphi}_i(\theta) \log(\tilde{\varphi}_i(\theta) / \hat{\varphi}_i^g(\theta)) d\theta;$$

i.e., for each  $i$ ,  $\hat{\varphi}_i^g(\theta)$  is asymptotically the “optimal” Gaussian approximation to  $\tilde{\varphi}_i(\theta)$ . No such relevant optimality holds for the product of densities, however: the (normalised) product of Gaussians, each of which is closest in the KL sense to  $\tilde{\varphi}_i(\theta)$ , is in general not the Gaussian closest to (the normalised version of)  $\prod \tilde{\varphi}_i(\theta)$ ; and indeed it may be very substantially different. In other words, as  $m \rightarrow \infty$ ,  $a$  and  $B$  do *not* in general minimise

$$\text{KL} \left( \left\{ \prod \tilde{\varphi}_i(\theta) / \int \left( \prod \tilde{\varphi}_i(\theta) \right) \right\} \parallel K(\theta, a, B) \right).$$

## 2.2 Kernel density estimate for $\hat{\varphi}_i(\theta)$

A second method we consider is to estimate each density  $\tilde{\varphi}_i(\theta)$  using a kernel density estimate (see for instance Silverman [1986] and Wand and Jones [1995]). A kernel density estimate based on Gaussian kernel functions (5) is

$$\hat{\varphi}_i^k(\theta) = \frac{1}{m} \sum_{j=1}^m K(\theta; \theta_{i(j)}^*, H_i), \quad (12)$$

where  $H_i$  is a bandwidth matrix. We follow the approach of Fukunaga [1972] in choosing the bandwidth matrix such that the shape of the kernel mimics the shape of the sample, in particular by taking  $H_i$  to be proportional to the sample covariance matrix,  $Q_i$ . Using bandwidth matrix

$$H_i = q \cdot m^{-2/(d+4)} Q_i, \quad (13)$$

where  $q > 0$  is a constant not dependent on  $m$ , ensures desirable behaviour as the sample size  $m \rightarrow \infty$ .

In particular, in terms of the little-o notation ( $a_m = o(b_m)$  as  $m \rightarrow \infty$  denotes  $\lim_{m \rightarrow \infty} |a_m/b_m| = 0$ ) and with  $E$  denoting expectation, using choice of bandwidth (13), subject to mild regularity conditions on  $\tilde{\varphi}_i(\theta)$ ,

$$E\{\hat{\varphi}_i^k(\theta)\} = \tilde{\varphi}_i(\theta) + o(1), \quad (14)$$

$$E\{\hat{\varphi}_i^k(\theta)^2\} = \tilde{\varphi}_i(\theta)^2 + o(1). \quad (15)$$

From (14)–(15), the bias,  $b\{\hat{\varphi}_i^k(\theta)\} = E\{\hat{\varphi}_i^k(\theta)\} - \tilde{\varphi}_i(\theta)$ , the variance,  $\text{var}\{\hat{\varphi}_i^k(\theta)\} = E\{\hat{\varphi}_i^k(\theta)^2\} - E\{\hat{\varphi}_i^k(\theta)\}^2$ , and the mean integrated squared error,

$$\text{MISE}\{\hat{\varphi}_i^k\} = E \int (\hat{\varphi}_i^k(\theta) - \tilde{\varphi}_i(\theta))^2 d\theta, \quad (16)$$

are all  $o(1)$ . These results generalise routinely to the case of a product of  $n$  kernel density estimates, that is, in which  $\prod \hat{\varphi}_i^k(\theta)$  is used as an estimator for  $\prod \tilde{\varphi}_i(\theta)$ . It follows that since the  $\theta_{i(j)}^*$  are independent for all  $i, j$ , then, using (14)–(15),

$$b\left\{\prod \hat{\varphi}_i^k(\theta)\right\} = \left\{\prod E \hat{\varphi}_i^k(\theta)\right\} - \prod \tilde{\varphi}_i(\theta) = o(1),$$

$$\text{var}\left\{\prod \hat{\varphi}_i^k(\theta)\right\} = \prod E\{\hat{\varphi}_i^k(\theta)^2\} - \prod \{E\hat{\varphi}_i^k(\theta)\}^2 = o(1).$$

$$\text{MISE}\left\{\prod \hat{\varphi}_i^k\right\} = E \int \left(\prod \hat{\varphi}_i^k(\theta) - \prod \tilde{\varphi}_i(\theta)\right)^2 d\theta = o(1).$$

Hence, in the sense defined by the latter equation, the density estimator  $\prod \hat{\varphi}_i^k(\theta)$  converges to the true density  $\prod \tilde{\varphi}_i(\theta)$  as  $m \rightarrow \infty$ .

Regarding the choice of  $q$  in (13), in certain settings it is possible to determine an optimal value. Suppose that the true density  $\tilde{\varphi}_i(\theta)$  is Gaussian and let  $\hat{\varphi}_i^k(\theta)$  in (12) be a kernel density estimate of  $\tilde{\varphi}_i(\theta)$ . Then

$$q = \{(d+2)/4\}^{-2/(d+4)} \quad (17)$$

is optimal in the sense that (13) is then an unbiased and consistent estimator of the bandwidth that minimises the leading term of the large- $m$  asymptotic

expansion of (16); see Wand and Jones [1995, p111]. Analogous calculations are rather more involved in the product case, however: even with the assumption that each  $\tilde{\varphi}_i(\theta)$  is Gaussian, no closed expression for  $q$  is possible. Hence, in the examples in the following section, §3, we opted to tune  $q$  in the heuristic way described by Wand and Jones [1995], starting with a large  $q$  (ten times that in (17)) then reducing it until “random” fluctuations begin to appear in the density estimates.

A consequence of using Gaussian kernel functions (5) in (12) is that the product of the density approximations is then itself a weighted mixture of  $(n-1)^m$  Gaussians,

$$\begin{aligned} \prod_{i=2}^n \hat{\varphi}_i^k(\theta) &= m^{(1-n)} \prod_{i=2}^n \sum_{j=1}^m K(\theta; \theta_{i(j)}^*, H_i) \\ &= m^{(1-n)} \sum_{j_2, \dots, j_n} \prod_{i=2}^n K(\theta; \theta_{i(j_i)}^*, H_i) \\ &= \sum_{j_2, \dots, j_n} w_{j_2, \dots, j_n} K(\theta; a_{j_2, \dots, j_n}, B_{j_2, \dots, j_n}), \end{aligned} \quad (18)$$

where expressions for the covariances  $B_{j_2, \dots, j_n}$ , means  $a_{j_2, \dots, j_n}$ , and weights  $w_{j_2, \dots, j_n}$ , analogous to those in (8)–(10), are given in the Appendix.

## 2.3 Estimating the posterior density

Sections §2.1 and §2.2 describe methods for computing the factor  $\prod \hat{\varphi}_i(\theta)$  in (3). For calculating an estimate of the full posterior,  $\hat{\pi}(\theta|\mathcal{X})$  in (3), we must multiply  $\prod \hat{\varphi}_i(\theta)$  by  $\pi(\theta)^{(2-n)}$  and normalise. Let us suppose that the prior is Gaussian,  $\pi(\theta) = K(\theta; \mu_{\text{pri}}, \Sigma_{\text{pri}})$ . For the case where we are using the Gaussian approximation,  $\hat{\varphi}_i^g(\theta)$  from (6), for each  $\hat{\varphi}_i(\theta)$ , then the posterior is

$$\hat{\pi}^g(\theta|\mathcal{X}) = K(\theta; \mu_{\text{post}}, \Sigma_{\text{post}}), \quad (19)$$

where

$$\Sigma_{\text{post}} = \left( (2-n)\Sigma_{\text{pri}}^{-1} + B^{-1} \right)^{-1}, \quad (20)$$

$$\mu_{\text{post}} = \Sigma_{\text{post}} \left( (2-n)\Sigma_{\text{pri}}^{-1} \mu_{\text{pri}} + B^{-1} a \right), \quad (21)$$

and  $a$  and  $B$  are as defined in (7).

If instead we use the kernel approximation,  $\hat{\varphi}_i^k(\theta)$  from (12), for each  $\hat{\varphi}_i(\theta)$ , then the posterior density is

$$\hat{\pi}^k(\theta|\mathcal{X}) = \frac{\sum_{j_2, \dots, j_n}^m w'_{j_2, \dots, j_n} K(\theta; a'_{j_2, \dots, j_n}, B'_{j_2, \dots, j_n})}{\sum_{j_2, \dots, j_n}^m w'_{j_2, \dots, j_n}}, \quad (22)$$

where expressions for  $B'_{j_2, \dots, j_n}$ ,  $a'_{j_2, \dots, j_n}$  and  $w'_{j_2, \dots, j_n}$  are in the Appendix.

## 2.4 Estimating the marginal likelihood

In some applications, especially when model comparison is of interest, it is useful to compute the marginal likelihood of the data given the model. The marginal likelihood is

$$\pi(\mathcal{X}) = \int \pi(\mathcal{X}|\theta) \pi(\theta) d\theta \quad (23)$$

$$= \left( \prod_{i=2}^n c_i \right) \int \left( \prod_{i=2}^n \varphi_i(\theta) \right) \pi(\theta)^{2-n} d\theta. \quad (24)$$

The unknown  $c_i$  can be estimated by  $\hat{c}_i = m/M_i$ , where  $M_i$  equals the number of ABC draws necessary in the  $i$ th interval to achieve  $m$  acceptances. For the integral in (24), using the Gaussian approximation

(7) leads to

$$\begin{aligned} & \int \left( \prod_{i=2}^n \hat{\varphi}_i^g(\theta) \right) \pi(\theta)^{2-n} d\theta \\ &= w \cdot (\det B)^{-1/2} \cdot (\det \Sigma_{\text{post}})^{1/2} \cdot (\det(2\pi \Sigma_{\text{pri}}))^{(n/2-1)} \times \\ & \exp \left\{ -\frac{1}{2} (a - \mu_{\text{pri}})^T ((2-n)^{-1} \Sigma_{\text{pri}} + B)^{-1} (a - \mu_{\text{pri}}) \right\}, \end{aligned} \quad (25)$$

whereas using the kernel approximation (12) gives

$$\int \left( \prod_{i=2}^n \hat{\varphi}_i^k(\theta) \right) \pi(\theta)^{2-n} d\theta = \sum_{j_2, \dots, j_n}^m w'_{j_2, \dots, j_n}. \quad (26)$$

## 2.5 Practical numerical calculations for the kernel approximation

Since expressions (18), (22), (26) for the kernel case involve sums with  $(n-1)^m$  terms, these expressions are largely of academic interest and are typically not suitable for practical calculations. For the examples in this paper we used a more direct numerical approach, first writing (4) as

$$g(\theta) = \exp \left( \sum_{i=2}^n h_i(\theta) \right) \pi(\theta),$$

where  $h_i(\theta) = \log(\varphi_i^k(\theta)/\pi(\theta))$ , and then evaluating  $h_i(\theta)$ ,  $\pi(\theta)$  and hence  $g(\theta)$  pointwise on a fine lattice. Performing calculations in this way on the log scale avoids underflow errors and improves numerical stability compared with trying to evaluate (4) directly. As a further check for robustness, we varied the lattice position and resolution to make sure the results were insensitive to the particular choices.

## 2.6 Sampling from the posterior distribution

In some circumstances it may be desirable to draw samples from the approximate posterior density. In the Gaussian case, drawing from (19) is straightforward. For the kernel case, (22), in principle sampling can be achieved by normalising the weights, randomly choosing a component with probability equal to these normalised weights, then sampling from the selected Gaussian component. But in practise, again, the large number of terms in (22) can preclude this approach. Other possibilities include using a Gibbs sampler, or sampling approximately using Gaussian mixtures with fewer components; see Sudderth et al. [2003].

## 3 Some examples

In this section we test PW-ABC on synthetic data from three models. The first, as a toy illustrative example, involves inferring from IID data the probability of success in a binomial experiment. Second, we consider an integer-valued time series model called INAR(1), a model for which the likelihood is available (albeit awkward to compute) and enables comparison of our approach with the “gold standard” MCMC approach. Third, we consider a stochastic Lotka–Volterra model, a simple example from a common class of models (which occur, for instance, in modelling stochastic chemical kinetics) in which the likelihood, and therefore many standard methods of inference, are unavailable. The datasets for each example are given in the Supplementary Material.

### 3.1 Binomial model

For this toy example we suppose the data is the set  $\mathcal{X} = \{x_1, \dots, x_{10}\}$  of  $n = 10$  observations from the model  $X_i \sim \text{Binom}(k_i = 100, p = 0.6)$ . We work in terms of the transformed parameter  $\theta = \text{logit}(p)$ , using a prior  $\pi(\theta) \sim N(0, 3^2)$ . For this model the data are IID, so that  $\pi(x_i | x_{i-1}, \theta) = \pi(x_i | \theta)$ . Exact samples from  $\varphi_i(\theta)$  can be obtained by sampling  $\theta^*$  from the prior, sampling  $X_i^* \sim \text{Binom}(100, \theta^*)$ , and then accepting  $\theta^*$  if and only if  $X_i^* = x_i$ . We follow the PW-ABC approach described in Section 2, drawing  $m = 5000$  samples from each  $\varphi_i(\theta)$ , using these samples to construct Gaussian  $\hat{\varphi}_i^g(\theta)$  and kernel density  $\hat{\varphi}_i^k(\theta)$  approximations, then using these density approximations to construct approximate posterior densities,  $\hat{\pi}^g(\theta | \mathcal{X})$  and  $\hat{\pi}^k(\theta | \mathcal{X})$ . Figure 1 shows that the approximate posterior densities are very close to the true posterior density for this example. The true log marginal likelihood,  $\log \pi(\mathcal{X})$ , computed by direct numerical integration of (23), is  $-31.39$ ; using approximation  $\hat{\varphi}_i^g(\theta)$  and (25) gives  $-31.44$ ; and using approximation  $\hat{\varphi}_i^k(\theta)$  and numerical integration of the left-hand side of (26) gives  $-31.48$ .

### 3.2 An integer-valued autoregressive model

Integer-valued time series arise in contexts such as modelling monthly traffic fatalities [Neal and Subba Rao, 2007] or the number of patients in a hospital at a sequence of time points [Moria et al., 2011]. Consider the following integer-valued autoregressive model of order  $p$ , known as INAR( $p$ ):

$$X_t = \sum_{i=1}^p \alpha_i \circ X_{t-i} + Z_t, \quad t \in \mathbb{Z}, \quad (27)$$

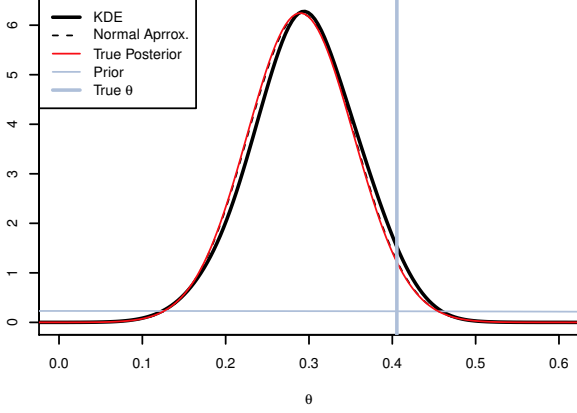


Figure 1: Results for the binomial model in §3.1. Shown are the true posterior density,  $\pi(\theta|\mathcal{X})$ , the posterior density approximations  $\hat{\pi}^g(\theta|\mathcal{X})$  and  $\hat{\pi}^k(\theta|\mathcal{X})$ , the prior, and the true  $\theta$ .

where  $Z_t$  for  $t > 1$  are independent and identically distributed integer-valued random variables with  $E[Z_t^2] < \infty$ , with the  $Z_t$  assumed to be independent of the  $X_t$ . Here we assume  $Z_t \sim Po(\lambda)$ . Each operator  $\alpha_i \circ$  denotes binomial thinning defined by

$$\alpha_i \circ W = \begin{cases} \text{Binomial}(W, \alpha_i), & W > 0, \\ 0, & W = 0, \end{cases} \quad (28)$$

for non-negative integer-valued random variable  $W$ . The operators  $\alpha_i \circ$ ,  $i = 1, \dots, p$ , are assumed to be independent.

We consider the simplest example of this model, INAR(1) [see, for example, Al-Osh and Alzaid, 1987], supposing that we have some observed data  $\mathcal{X} = \{x_1, \dots, x_n\}$  from this model and wish to make inference for the parameters  $(\alpha, \lambda)$ . We generated  $n = 100$  observations from an INAR(1) process using parameters  $(\alpha, \lambda) = (0.7, 1)$  and  $X(0) = 10$ ;

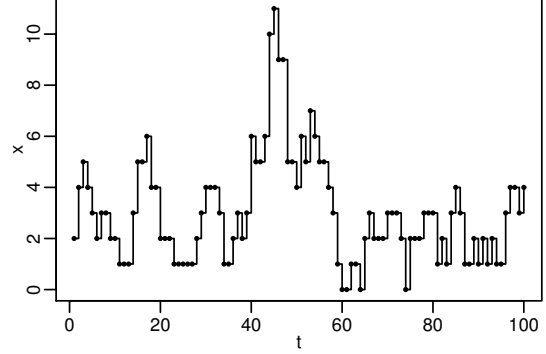


Figure 2: The realisation of an INAR(1) process used in the example of §3.2, of length  $n = 100$ , generated using  $\alpha = 0.7$  and  $\lambda = 1.0$ .

the realisation is plotted in Figure 2. Working in terms of the transformed parameter,  $\theta = (\theta_1, \theta_2) = (\text{logit}(\alpha), \log(\lambda))$ , we used a prior of  $\text{Norm}(0, 3^2)$  for each of  $\theta_1$  and  $\theta_2$ . For the EBC algorithm, the probability of acceptance is around  $10^{-100}$ , which is prohibitively small; even the ABC algorithm requires a value of  $\varepsilon$  so large that sequential approaches (e.g., SMC-ABC) are needed.

Using PW-ABC with  $s(\cdot) = \text{Identity}(\cdot)$  and  $\varepsilon = 0$  we were able to draw exact samples from  $\varphi_i(\theta)$  for all of the  $i = 2, \dots, 100$  factors, and still achieve acceptance rates of around 9%, on average. Figure 3 shows an estimate of the posterior density,  $\pi(\theta|\mathcal{X})$  based on a gold-standard MCMC approach, together with Gaussian- and kernel-based PW-ABC approximations,  $\hat{\pi}^g(\theta|\mathcal{X})$  and  $\hat{\pi}^k(\theta|\mathcal{X})$ , with  $m = 10,000$  samples for each  $\varphi_i(\theta)$ . The Figure shows good agreement between the MCMC posterior and the kernel approximation,  $\hat{\pi}^k(\theta|\mathcal{X})$ , but poor agreement with the Gaussian approximation  $\hat{\pi}^g(\theta|\mathcal{X})$ . The poor performance of  $\hat{\pi}^g(\theta|\mathcal{X})$  is probably caused by several

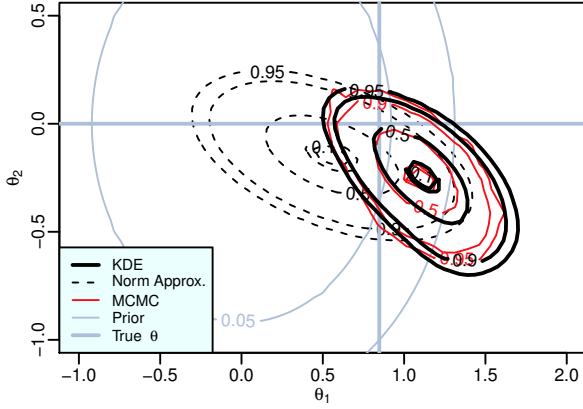


Figure 3: Results for the INAR(1) example of §3.2. Shown are an MCMC approximation to the posterior density,  $\pi(\theta|\mathcal{X})$ , the posterior density approximations  $\hat{\pi}^g(\theta|\mathcal{X})$  and  $\hat{\pi}^k(\theta|\mathcal{X})$ , the prior, and the true  $\theta$ . The numbers on the contours denote the probability mass that they contain.

of the densities  $\varphi_i(\theta)$  being substantially different from Gaussian; see Figure 4 which shows  $\hat{\varphi}_{50}^g(\theta)$  and  $\hat{\varphi}_{50}^k(\theta)$ —the latter suggests that the density is far from Gaussian. Using Gaussian approximations to non-Gaussian  $\varphi_i(\theta)$  appears to have a strong impact on the accuracy of approximation  $\hat{\pi}^g(\theta|\mathcal{X})$ , even as in this case where the true posterior  $\pi(\theta|\mathcal{X})$  is reasonably close to a Gaussian (cf. Fig. 3).

For this example, estimates of the marginal likelihood,  $\log \pi(\mathcal{X})$ , are as follows: by direct numerical integration of (23),  $-161.1$ ; using approximation  $\hat{\varphi}_i^g(\theta)$  and (25),  $-185.7$ ; and by using  $\hat{\varphi}_i^k(\theta)$  and numerical integration of the left-hand side of (26),  $-163.2$ .

We have used  $p = 1$  for this example so that the likelihood is available, enabling comparison with MCMC and calculation of the true marginal likelihood. However, we stress that PW-

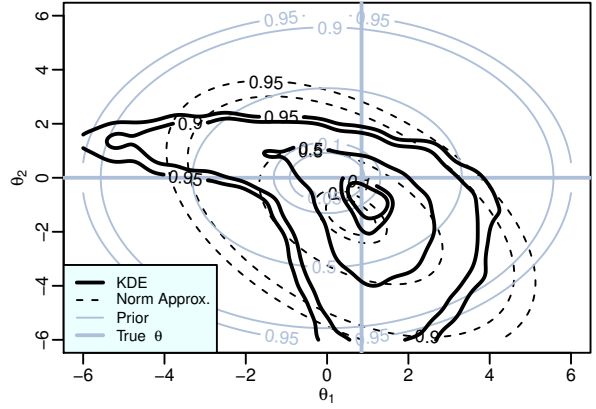


Figure 4: For the INAR(1) example, an example of a factor with a “non-Gaussian” density: here  $\hat{\varphi}_{50}^g(\theta)$  and  $\hat{\varphi}_{50}^k(\theta)$  are substantially different from each other.

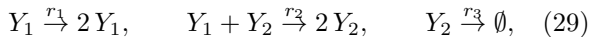
ABC can be applied equally easily for  $p > 1$ , a case for which the likelihood is essentially intractable and therefore one has to resort to either exact but less direct methods (such the Expectation–Maximization (EM) algorithm or data-augmented MCMC, both of which involve treating the terms  $\alpha_i \circ X_{t-i}$  and  $Z_t$  as missing data) or methods of approximate inference, such as conditional least squares which involves minimizing  $\sum_t (X_t - E[X_t|X_{t-1}])^2$ ; see, for example, McKenzie [2003] and references therein.

### 3.3 Stochastic Lotka–Volterra Dynamics

The stochastic Lotka–Volterra (LV) model is a model of predator–prey dynamics and an example of a stochastic discrete-state-space continuous-time Markov process [see, for example, Wilkinson, 2011]. Models of this type

commonly arise when modelling chemical kinetics. Predator-prey dynamics can be thought of in chemical kinetics terms: the predators and prey are two populations of “reactants” subject to three “reactions”, namely prey birth, predation and predator death. Exact simulation of such models is straightforward, e.g., using the algorithm of Gillespie (1977). Inference is simple if the type and precise time of each reaction is observed. However, a more common setting is where the population sizes are only observed at discrete time points. In this case the number of reactions that have taken place is unknown and therefore the likelihood is not available and hence inference is much more difficult. Reversible-jump MCMC has been developed in this context [Boys et al., 2008] but it requires substantial expertise and input from the user to implement. Particle MCMC (pMCMC) methods [Andrieu et al., 2010], which provide an approximation to the likelihood via a Sequential Monte Carlo (SMC) algorithm within an MCMC algorithm, have recently been proposed for stochastic chemical kinetics models [Golightly and Wilkinson, 2011]. Although being computationally intensive, such methods can work reliably provided the process is observed with measurement error. The R package `smfbsb`, which accompanies Wilkinson [2011], contains a pMCMC implementation designed for stochastic chemical kinetics models, and we use this package to compare results for PW-ABC and pMCMC for the following example.

Let  $Y_1$  and  $Y_2$  denote the number of prey and predators respectively, and suppose  $Y_1$  and  $Y_2$  are subject to the following reactions



which respectively represent prey birth, predation and predator death. We consider the problem of making inference about the rates  $(r_1, r_2, r_3)$  based on observations of  $Y_1$  and  $Y_2$  made at fixed intervals.

We generated a realisation from the stochastic LV example of Wilkinson [2011, page 208], that is, model (29) using  $(r_1, r_2, r_3) = (1, 0.005, 0.6)$ ,  $Y_1(0) = 50$  and  $Y_2(0) = 100$ . We performed inference in terms of transformed parameters,  $\theta = (\theta_1, \theta_2, \theta_3) = (\log r_1, \log r_2, \log r_3)$ , this time with priors  $\pi(\theta_1) \sim \text{Norm}(\log(0.7), 0.5)$ ,  $\pi(\theta_2) \sim \text{Norm}(\log(0.005), 0.5)$ , and  $\pi(\theta_3) \sim \text{Norm}(\log(0.3), 0.5)$ . We again applied PW-ABC using  $s(\cdot) = \text{Identity}(\cdot)$  and  $\varepsilon = 0$ , in other words requiring an exact match between the observed and the simulated observations, to draw samples of size  $m = 10,000$  for each  $\varphi_i(\theta)$ .

To obtain pMCMC results we found it necessary to assume an error model for the observations, hence we assumed errors to be IID Gaussian with mean zero and standard deviation equal to 2. Results are displayed in Figure 3.3, which shows plots for univariate and pairwise bivariate marginal posterior densities for the pMCMC results, and for the PW-ABC approximations,  $\hat{\pi}^g(\theta|\mathcal{X})$  and  $\hat{\pi}^k(\theta|\mathcal{X})$ . Both of the PW-ABC approximations agree well with each other and with the pMCMC results for this example.

## 4 Conclusion and Discussion

PW-ABC aims at using Markov structure to lower the computational cost involved in ABC. The major advantage of PW-ABC over ABC is that for a given summary statistic,  $s(\cdot)$ , and tolerance,  $\varepsilon$ , acceptance rates for sampling from  $\varphi_i(\theta)$  are overwhelmingly higher in comparison to sampling from  $\pi(\theta|\mathcal{X})$ . Po-

tentially PW-ABC enables the choice  $s = \text{Identity}(\cdot)$ , sidestepping the difficult task of deciding what summary statistic is appropriate (i.e., “approximately sufficient”) and eliminating ensuing biases that result from the choice.

Having sampled from each  $\varphi_i(\theta)$  the question then becomes how to estimate  $\pi(\theta|\mathcal{X})$  using these samples. PW-ABC works by constructing density approximations  $\hat{\varphi}_i(\theta)$  to each  $\varphi_i(\theta)$ . The approach of taking  $\hat{\varphi}_i(\theta)$  to be Gaussian, with moments matched to the sample moments, is computationally cheap, and if the prior is also taken to be Gaussian then there is a closed form expression for the Gaussian posterior density and marginal likelihood, making calculations extremely fast. Taking  $\hat{\varphi}_i(\theta)$  to be Gaussian is perhaps adequate in many applications: performance was strong in two of the three examples we considered. The poor performance in the INAR example of §3.2 is probably because some of the  $\varphi_i(\theta)$  are so different from Gaussian. It is striking to see an effect so strong when the true posterior is so close to Gaussian. Unfortunately, increasing the number,  $m$ , of ABC samples is no remedy to this problem: as  $m \rightarrow \infty$ , the product of Gaussians, itself a Gaussian, in general does not converge to the Gaussian closest in the Kullback–Leibler sense to the target density.

In terms of asymptotic performance, using the kernel approximation,  $\hat{\varphi}_i^k(\theta)$ , for  $\hat{\varphi}_i(\theta)$  is preferable since, in this case, the estimated posterior density converges to the target as  $m \rightarrow \infty$ . The kernel approach is computationally more demanding, however, and also calls for a heuristic choice of a scalar smoothing parameter. In spite of its asymptotic properties, its practical use is probably limited to problems in which  $\theta$  has small dimension.

A more general possibility which we will explore in future work is to let  $\hat{\varphi}_i(\theta)$  be a mixture of, say,  $u$  Gaussians. This encompasses (6) and (12) as special cases (with  $u = 1$  and  $u = m$  respectively). For a general mixture model for  $\hat{\varphi}_i(\theta)$ , each of the component Gaussians is parameterised by a scalar weight, a mean vector and a covariance matrix which need to be determined. We would envisage regularising, e.g., by setting each covariance to be equal up to scalar multiplication, perhaps as for (12) taking the covariance equal to the sample covariance, and then fitting each  $\hat{\varphi}_i(\theta)$  based on the samples from  $\varphi_i(\theta)$  using, say, an EM algorithm. This approach is a compromise between (6) and (12). It does not share the property of (12) that estimated densities converge to the true densities as  $m \rightarrow \infty$ , but on the other hand it is computationally much less involved and offers much extra freedom and flexibility over (6), particularly for dealing with multimodal densities. If  $u$  is taken sufficiently small then it may be feasible to work explicitly with the  $(n-1)^u$ -term resulting Gaussian mixture,  $\prod \hat{\varphi}_i(\theta)$ , enabling explicit calculations involving the posterior density, such as computing the marginal likelihood, analogous to (25), and direct sampling from the approximate posterior density (see §2.6).

Several further generalisations of the PW-ABC approach are possible. In (1), each of the  $n-1$  factors  $\pi(x_i|x_{i-1}, \theta)$ ,  $i = 2, \dots, n$  is the likelihood for a single data point conditional on the previous. An alternative possibility is to factorise the likelihood into fewer factors, with each corresponding to a “block” of multiple observations, e.g.,  $\pi(x_{i+v_i}, x_{i+v_i-1}, \dots, x_i|x_{i-1}, \theta)$  for some choice of  $v_i$ , and the factorised likelihood becomes a product over

the relevant subset of  $i = 2, \dots, n$ . To an extent this potentially reintroduces difficulties that with PW-ABC we sought to avoid, namely lower acceptance rates leading to a possible need to use a summary statistic and non-zero tolerance (and the ensuing ABC error they bring). On the other hand, we might expect, owing to the central limit theorem, that a factor  $\varphi_i(\theta)$  which depends on several data points will be closer to Gaussian than a factor dependent on only a single data point, and hence that (6) and (12) (especially the former) will perform better.

If using larger “blocks” of data in the factorisation makes it necessary to use a non-zero tolerance  $\varepsilon > 0$  (or if  $\varepsilon > 0$  is necessary even when using a single observation per factor) then there are theoretical advantages to using what Fearnhead and Prangle [2012] call “noisy ABC”. In the context of this paper, noisy ABC would involve replacing the summary statistic  $s(\cdot)$  with a random version  $s'(\cdot) = s(\cdot) + \alpha$ , where  $\alpha$  is a realisation of a random variable with density uniform on the ball  $\{x | d(s(\mathcal{X}), x) \leq \varepsilon\}$ . Using noisy ABC ensures that, under mild regularity conditions, as  $n \rightarrow \infty$ , the posterior converges to a point mass at the true parameter value; see §2.2 of Fearnhead and Prangle [2012].

Recently, we have learnt of an interesting paper by Barthelmé and Chopin [2011] who have developed an approach termed *Expectation Propagation-ABC* (EP-ABC) that shares similarities with ours. EP-ABC is an ABC adaptation of the Expectation Propagation approach developed by Minka [2001]. EP-ABC uses essentially the same factorisation as (2) and makes a Gaussian approximation to the density of each factor analogous to (6). But then EP-ABC proceeds rather differently: instead of drawing ABC samples for, say,

the  $i$ th factor by sampling from the prior, EP-ABC draws samples from an iteratively updated pseudo-prior. The pseudo-prior is a Gaussian approximation to the component of the posterior that involves all the data *except* those pertaining to the  $i$ th factor. The use of the pseudo-prior offers a high acceptance rate in the ABC sampling and so EP-ABC can potentially lead to an extremely fast approximation to the full posterior  $\pi(\theta | \mathcal{X})$ . A disadvantage is that conditions sufficient for the convergence of EP-ABC (or even the simpler deterministic EP) are not known. Also, as with using PW-ABC with (7), since EP-ABC uses a Gaussian approximation for each factor, it is potentially ill-suited to problems with complicated (e.g. multimodal or otherwise non-Gaussian) likelihoods; convergence of the product density is not assured to any “optimal” approximation to the target posterior. A promising direction for future work will be to investigate adapting the EP-ABC idea of sampling from a pseudo-prior to the ideas in this paper of using kernel (or Gaussian mixture) density estimates for each likelihood factor.

## Acknowledgements

S.R. White was supported by the (UK) Medical Research Council [Unit Programme number U105260794] and the EPSRC [University of Nottingham, Bridging the Gaps]. The authors are grateful to Andy Wood for useful suggestions about the manuscript.

## Appendix

Expression for  $B_{j_2, \dots, j_n}$ ,  $a_{j_2, \dots, j_n}$ , and  $w_{j_2, \dots, j_n}$  in (18), analogous to (8)–(10), are as follows:

$$B_{j_2, \dots, j_n} = \left( \sum_{i=2}^n H_i^{-1} \right)^{-1},$$

$$a_{j_2, \dots, j_n} = B_{j_2, \dots, j_n} \left( \sum_{i=2}^n H_i^{-1} \theta_{i(j_i)}^* \right),$$

$$w_{j_2, \dots, j_n} = m^{(1-n)} \det(2\pi B_{j_2, \dots, j_n})^{1/2} \prod_{i=2}^n \det(2\pi H_i)^{-1/2} \times$$

$$\prod_{s=2}^n \prod_{t>s}^n \exp \left( -\frac{1}{2} (\theta_{s(j_s)}^* - \theta_{t(j_t)}^*)^T R_{st} (\theta_{s(j_s)}^* - \theta_{t(j_t)}^*) \right),$$

$$R_{st} = H_s^{-1} B_{j_2, \dots, j_n} H_t^{-1}.$$

Expressions for  $B'_{j_2, \dots, j_n}$ ,  $a'_{j_2, \dots, j_n}$ , and  $w'_{j_2, \dots, j_n}$  in (22) are given respectively by the right-hand sides of (20), (21), and (25) with  $B$  replaced by  $B_{j_2, \dots, j_n}$ ,  $a$  replaced by  $a_{j_2, \dots, j_n}$ , and  $w$  replaced by  $w_{j_2, \dots, j_n}$ .

## References

- M.A. Al-Osh and A.A. Alzaid. First-order integer-valued autoregressive (INAR(1)) process. *J. Time Ser. Anal.*, 8(3):261–275, 1987.
- C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(3):269–342, 2010.
- S. Barthelmé and N. Chopin. Expectation-Propagation for Summary-Less, Likelihood-Free Inference. *ArXiv e-prints*, 2011.
- M.A. Beaumont, W. Zhang, and D.J. Balding. Approximate Bayesian Computation in Population Genetics. *Genetics*, 162(4):2025–2035, December 2002.
- M.G.B. Blum and O. Franois. Non-linear regression models for approximate bayesian computation. *Statistics and Computing*, 20:63–73, 2010.
- R.J. Boys, D.J. Wilkinson, and T.B. Kirkwood. Bayesian inference for a discretely observed stochastic kinetic model. *Stat. Comput.*, 18(2):125–135, 2008.
- T. A. Dean, S. S. Singh, A. Jasra, and G. W. Peters. Parameter Estimation for Hidden Markov Models with Intractable Likelihoods. *ArXiv e-prints*, 2011.
- P. Fearnhead and D. Prangle. Constructing summary statistics for approximate bayesian computation: semi-automatic approximate bayesian computation. *RSS Series B*, 2012. In Press.
- K. Fukunaga. *Introduction to statistical pattern recognition*. Electrical science series. Academic Press, 1972.
- E. Gabriel, D.J. Wilson, A.J. Leatherbarrow, J. Cheesbrough, S. Gee, E. Bolton, A. Fox, P. Fearnhead, C.A. Hart, and P.J. Diggle. Spatio-temporal epidemiology of campylobacter jejuni enteritis, in an area of northwest england, 2000–2002. *Epidemiology and Infection*, 138:1384–1390, 2010.
- A. Golightly and D.J. Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle markov chain monte carlo. *Interface Focus*, 1(6):807–820, 2011.
- J.M. Marin, P. Pudlo, C. Robert, and R. Ryder. Approximate bayesian computational methods. *Statistics and Computing*, pages 1–14, 2011. doi: 10.1007/s11222-011-9288-2.

- P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain Monte Carlo without likelihoods. *Proc. Natl Acad. Sci. USA*, 100(26):15324, 2003.
- E. McKenzie. Discrete variate time series. In *Stochastic processes: modelling and simulation*, volume 21 of *Handbook of Statist.*, pages 573–606. North-Holland, Amsterdam, 2003.
- T. McKinley, A. Cook, and R. Deardon. Inference in epidemic models without likelihoods. *Intl. J. Biostat.*, 5:24, 2009.
- T.P. Minka. Expectation propagation for approximate bayesian inference. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, UAI '01, pages 362–369, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1-55860-800-1. URL <http://dl.acm.org/citation.cfm?id=647235.720257>.
- D. Moria, P. Puig, J. Ros, A. Vilella, and A. Trilla. A statistical model for hospital admissions caused by seasonal diseases. *Statistics in Medicine*, 30(26):3125–3136, 2011.
- P. Neal and T. Subba Rao. MCMC for integer-valued ARMA processes. *J. Time Ser. Anal.*, 28(1):92–110, 2007.
- J.K. Pritchard, M.T. Seielstad, A. Perez-Lezaun, and M.W. Feldman. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Mol. Biol. Evol.*, 16(12):1791–1798, 1999.
- B.W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, London - New York, 1986.
- E.B. Sudderth, A.T. Ihler, W.T. Freeman, and A.S. Willsky. Nonparametric belief propagation. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 1:605, 2003.
- T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M.P. Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J. R. Soc. Interface*, 6:187–202, 2009.
- N.G. Van Kampen. *Stochastic Processes in Physics and Chemistry*. North-Holland Personal Library. Elsevier, 2007.
- P. Wand and C. Jones. *Kernel Smoothing*. Monographs on Statistics and Applied Probability. Chapman & Hall, 1995.
- D.J. Wilkinson. *Stochastic modelling for systems biology*. Chapman & Hall/CRC, 2011. 2nd Edition.

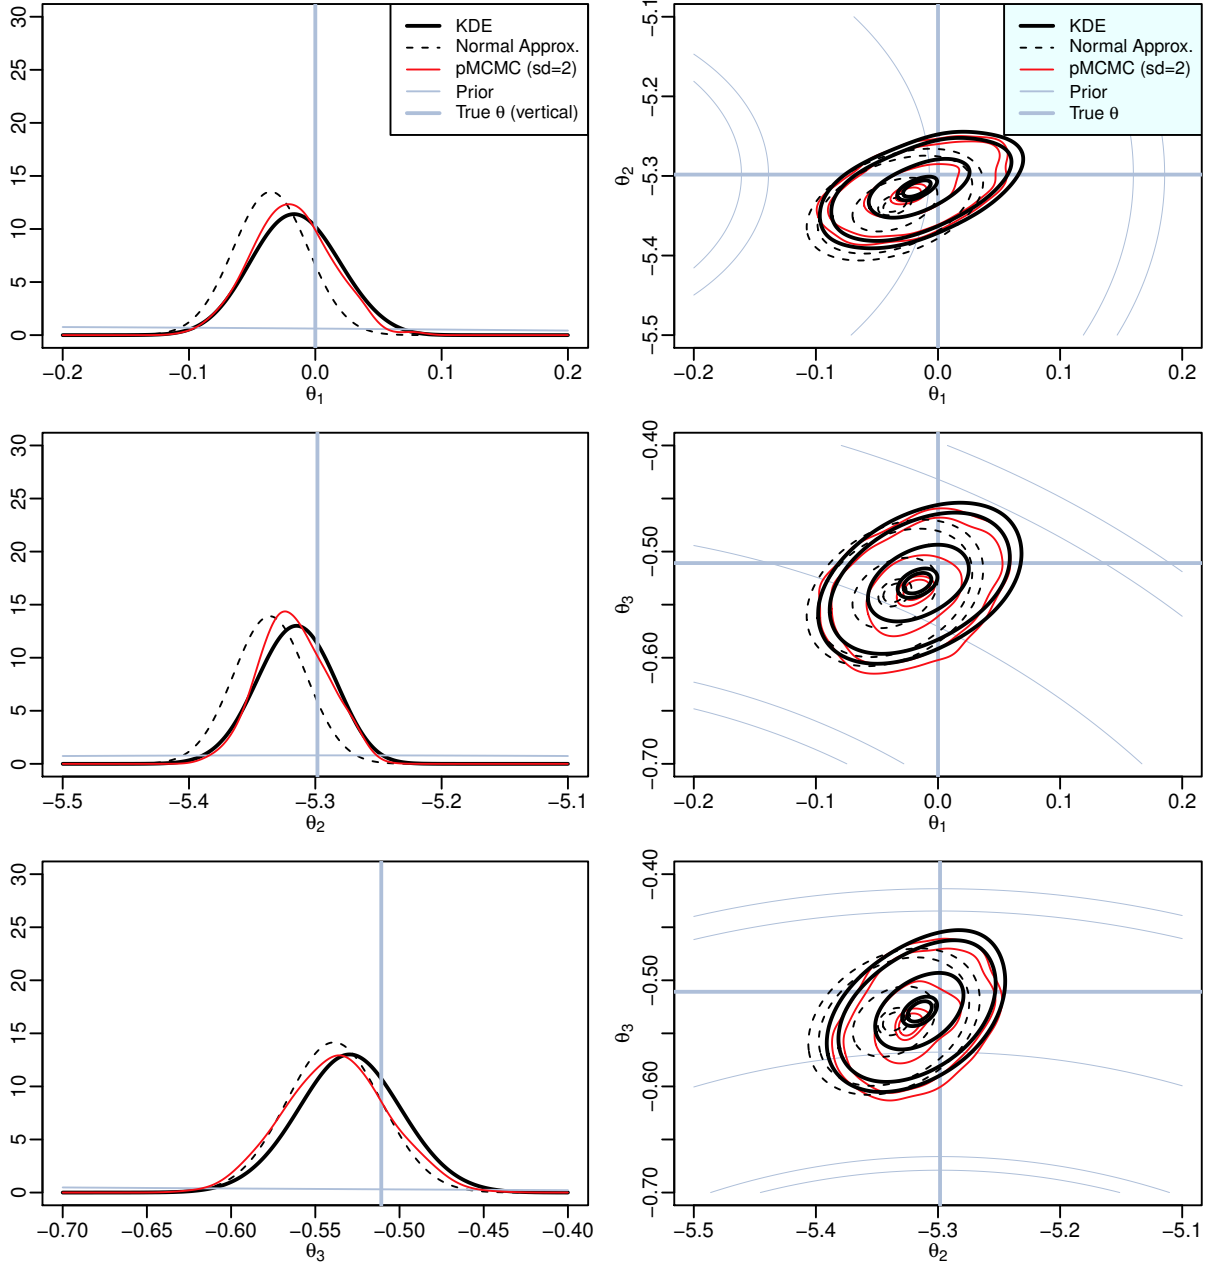


Figure 5: Results for the Lotka–Volterra example of §3.3, showing univariate and bivariate marginal posterior densities of  $\theta$  based on a posterior sample from a pMCMC algorithm, and from the Gaussian- and kernel-based PW-ABC approximations,  $\hat{\pi}^g(\theta|\mathcal{X})$  and  $\hat{\pi}^k(\theta|\mathcal{X})$ . For the kernel approximation we used  $q = 5$  as the smoothing parameter in (13). The contours shown in the bivariate plots are those that contain 5%, 10%, 50%, 90% and 95% of probability mass.